

09-15-00

A




IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of: Abraham R. Matthews et al.

Title: SYSTEM AND METHOD FOR ALLOCATING NETWORKING RESOURCES

Attorney Docket No.: 1384.003US1



PATENT APPLICATION TRANSMITTAL

BOX PATENT APPLICATION

Commissioner for Patents
Washington, D.C. 20231

We are transmitting herewith the following attached items and information (as indicated with an "X"):

- ☒ Return postcard.
- ☒ Utility Patent Application under 37 CFR § 1.53(b) comprising:
 - ☒ Specification (27 pgs, including claims numbered 1 through 1 and a 1 page Abstract).
 - ☒ Formal Drawing(s) (20 sheets).
 - ☒ Unsigned Combined Declaration and Power of Attorney (4 pgs).
 - ☒ Appendix A (2 pgs.)
 - ☒ Appendix B (7 pgs.)

The filing fee (NOT ENCLOSED) will be calculated as follows:

	No. Filed	No. Extra	Rate	Fee
TOTAL CLAIMS	1 - 20 =	0	x 18 =	\$0.00
INDEPENDENT CLAIMS	1 - 3 =	0	x 78 =	\$0.00
MULTIPLE DEPENDENT CLAIMS PRESENTED				\$0.00
BASIC FEE				\$690.00
TOTAL				\$690.00

THE FILING FEE WILL BE PAID UPON RECEIPT OF THE NOTICE TO FILE MISSING PARTS.

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
P.O. Box 2938, Minneapolis, MN 55402 (612-373-6900)

By: Thomas F. Brennan
Atty: Thomas F. Brennan
Reg. No. 35,075

Customer Number 21186

"Express Mail" mailing label number: EL618477168US

Date of Deposit: September 13, 2000

This paper or fee is being deposited on the date indicated above with the United States Postal Service pursuant to 37 CFR 1.10, and is addressed to the Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

UNITED STATES PATENT APPLICATION

**SYSTEM AND METHOD FOR
ALLOCATING NETWORKING RESOURCES**

INVENTORS

Abraham Rabindranath Matthews

Anna Berenberg

Kevin Lin

Schwegman, Lundberg, Woessner, & Kluth, P.A.

1600 TCF Tower

121 South Eighth Street

Minneapolis, Minnesota 55402

ATTORNEY DOCKET 1384.003US1

SYSTEM AND METHOD FOR ALLOCATING NETWORKING RESOURCES

Field of the Invention

The present invention is related to networking systems, and more particularly to a system and method for allocating networking resources within a wide area network (WAN).

Background Information

Internet or WAN service providers are operating in a crowded marketplace where cost effectiveness is critical. Operational costs present a significant challenge to service providers. Cumbersome, manual provisioning processes are the primary culprits. Customer orders must be manually entered and processed through numerous antiquated back-end systems that have been pieced together. Once the order has been processed, a truck roll is required for onsite installation and configuration of Customer Premises Equipment (CPE), as well as subsequent troubleshooting tasks.

Presently, the delivery of firewall services requires the deployment of a specialized piece of Customer Premises Equipment (CPE) to every network to be protected. This model of service delivery creates an expensive up-front capital investment, as well as significant operational expenses that are associated with onsite installation and management of thousands of distributed devices. The results are service delivery delays, increased customer start-up costs and/or thinner service provider margins.

The slow and expensive process of deploying firewall services cuts into margins and forces significant up-front charges to be imposed on the customer. In order to be successful in today's market, service providers must leverage the public network to offer high-value, differentiated services that maximize margins while controlling capital and operational costs. These services must be rapidly provisioned and centrally

managed so that time-to-market and, more importantly, time-to-revenue are minimized. Traditional methods of data network service creation, deployment, and management present significant challenges to accomplishing these goals, calling for a new network service model to be implemented.

Enterprise customers are increasingly demanding cost-effective, outsourced connectivity and security services, such as Virtual Private Networks (VPNs) and managed firewall services. Enterprise networks are no longer segregated from the outside world; IT managers are facing mounting pressure to connect disparate business units, satellite sites, business partners, and suppliers to their corporate network, and then to the Internet. This raises a multitude of security concerns that are often beyond the core competencies of enterprise IT departments. To compound the problem, skilled IT talent is an extremely scarce resource. Service providers, with expert staff and world-class technology and facilities, are well positioned to deliver these services to enterprise customers.

What is needed is a system and method for providing managed network services that are customizable for each customer's need.

Description of the Preferred Embodiments

In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the ways used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these

quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar computing device, that manipulates and transforms data represented as physical (e.g., electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

While IT managers clearly see the value in utilizing managed network services, there are still barriers to adoption. Perhaps the most significant of these is the fear of losing control of the network to the service provider. In order to ease this fear, a successful managed network service offering must provide comprehensive visibility to the customer, enabling them to view configurations and performance statistics, as well as to request updates and changes. By providing IT managers with powerful Customer Network Management (CNM) tools, one can bolster confidence in the managed network service provider and can actually streamline the service provisioning and maintenance cycle.

While service providers recognize the tremendous revenue potential of managed firewall services, the cost of deploying, managing and maintaining such services via traditional CPE-based methods is somewhat daunting. Service providers are now seeking new service delivery mechanisms that minimize capital and operational costs while enabling high-margin, value-added public network services that are easily provisioned, managed, and repeated. Rolling out a network-based managed firewall service is a promising means by which to accomplish this. Deploying an IP Service

Delivery Platform in the service provider network brings the intelligence of a managed firewall service out of the customer premises and into the service provider's realm of control.

One such IP Service Delivery Platform 10 is shown in Fig. 1. In the embodiment shown in Fig. 1, IP Service Delivery Platform 10 includes three distinct components: an intelligent, highly scalable IP Service Processing Switch 12, a comprehensive Service Management System (SMS) 14 and a powerful Customer Network Management (CNM) system 16. Service Management System (SMS) 14 is used to enable rapid service provisioning and centralized system management. Customer Network Management (CNM) system 16 provides enterprise customers with detailed network and service performance systems, enable self-provisioning. At the same time, system 16 eases IT managers fears of losing control of managed network services.

In one embodiment, such as is shown in Fig. 2 for a network-based managed firewall service model, the service provider replaces the high-capacity access concentration router at the POP with an IP Service Processing Switch 12. This is a higher-capacity, more robust, and more intelligent access switch, with scalable processing up to 100+ RISC CPUs. Just as with the access router, additional customer access capacity is added via installing additional port access blades to the IP Service Processing Switch chassis. Unlike conventional access routers, however, additional processor blades can be added to switch 12 to ensure wire-speed performance and service processing.

The intelligence resident in IP Service Processing Switch 12 eliminates the need to deploy CPE devices at each protected customer site. Deployment, configuration, and management of the managed firewall service all take place between IP Service Processing Switch 12 and its Service Management System 14. In the embodiment shown in Fig. 2, Service Management System 14 resides on a high-end UNIX platform at the service provider NOC.

In one embodiment, the customer has the ability to initiate service provisioning and augmentation via a web-based Customer Network Management tool residing, e.g.,

at the customer's headquarters site. This is an entirely different service delivery paradigm, requiring little or no truck rolls and little or no on-site intervention.

In one embodiment, switch 12 is a 26-slot services processing switch that marries scalable switching, routing and computing resources with an open software architecture to deliver computationally-intense IP services such as VPNs with scalable high performance. In one embodiment, switch 12 has a high-speed 22 Gbps redundant dual counter-rotating ring midplane. Slots are configured with four types of Service Blades: Control, Access, Trunk and Processor blades with specialized processing which enables a range of high-performance services including route forwarding, encryption and firewalls.

Service providers can use switch 12's virtual routing capabilities, and its ability to turn IP services into discrete and customized objects, to segment and layer services for the first time for tens of thousands of discrete subscriber corporations. In addition, processor capacity can be added to switch 12 by adding new processor blades.

In one embodiment switch 12 includes an operating system which dynamically distributes services to switch 12 processors.

In one embodiment, the 26-slot services processing switch corrects for failures using the redundant counter-rotating ring midplane.

In one embodiment, each Service Blade automatically fails-over to a backup.

One embodiment of a switch 12 is described in "System and Apparatus for Delivering Security Services," filed herewith, the description is incorporated herein by reference.

In one embodiment, switch 12 is designed to integrate seamlessly into a SP's preexisting network, whether that be through support of open routing protocols or through its Frame Relay to IPSec interworking solution that integrates new IP-based networks into a corporation's preexisting Frame Relay cloud.

The operating system will be described next.

In one embodiment, switch 12 includes a network operating system (NOS) 20. In one embodiment, network operating system 20 enables switch 12 to create discrete customized services to specific subscriber corporations by providing them each with a

different configuration of service object groups. NOS 20 enables objects within these object groups to be distributed dynamically to customized processors so that application services are receiving the right level of computational support.

In one embodiment, NOS 20 is designed as an open Application Program Interface (API) that allows general-purpose software or new advanced IP services to be ported into the platform from best of breed third parties in a continual fashion, helping to enrich service provider' investment over time.

In one embodiment, NOS 20 includes a distributed messaging layer (DML) 22 component, an object manager (OM) component 24 layered over DML, control blade redundancy (CBR) 26 for redundant system controllers, a resource manager (RM) 28 for managing separate resource elements and a resource location service (RLS) 30. Resource location service 30 provides load balancing across capable processor elements (Pes) to create an object. PE selection is based on predefined constraints.

In one embodiment, CBR 26 is layered over DML 22 and OM 24 as shown in Fig. 3.

In one embodiment, Object Manager 24 consists of three layers as shown on Fig. 4. The upper layer titled *OM Controller and Database* (OMCD) 40 is concerned with managing the VPN and VR configuration. This is the agent that deals with the configuration manager directly. Middle layer 42 entitled *OM Object Routing and Interface Global* is concerned with managing global (across the switch system) object groups and object configurations. Lower layer 44 entitled *OM Object Routing and Interface* (OMORI) is concerned with managing local objects and groups as well as routing control information between address spaces based on the location of objects, and interfacing with the object via method invocation.

In one embodiment, the IPSX object database consists of two types of databases: Global (managed on Master Control Blade by OMORIG) and distributed local databases (managed by OMORI agents on every PE present in the system). In one such embodiment, the global database is a superset of the extracts from local databases.

One such network operating system 20 is described in "Switch Management System and Method," filed herewith, the description of which is incorporated herein by reference.

In one embodiment, objects represent a basic unit of management for purposes of fault tolerance, computational load balancing, etc. One or more adjacent protocol modules can be placed into a single object. In addition, it is also possible that a module can be split across two objects.

In one embodiment, virtual routers (VRs) are mapped onto the physical processor elements of switch 12. In one such embodiment, each VR is logical mirror of a stand-alone hardware router. The virtual router (VR) construct is at the heart of the system 10 service segmentation and layering method. By enabling each subscriber entity to have a pool of VRs across multiple POPs for its VPN, the system is able to support tens of thousands of discrete entities within a single platform.

Just like a traditional router, a VR has its own route-forwarding table. In addition, each VR manages a defined set of routing protocols and supports a number of interfaces. In one embodiment, VRs support not only physical interfaces, but also Virtual Interfaces (VIs). This enables communication between VRs within the same switch 12.

In one embodiment, each switch 12 can support up to 10,000 VRs per switch.

Fundamentally, the system VR (ISPX VR) is no different than any other VR in its make-up. It differs from the other VRs, however, in that it is the router that aggregates traffic from all other VRs and forwards it out to another switch 12 or out to the Internet. By isolating the subscriber level VRs from a direct Internet adjacency, system 10 is able to minimize its complexity to one element in the system. In addition, the ISPX VR performs routing protocol isolation, for instance by interacting with the core via BGP-4, while running OSPF or RIP to a subscriber VR. It also allows the system to optimize resources, by keeping the subscriber VR routing tables small.

A virtual router performs functions similar to a real/normal router. One or more VR would be setup for a subscriber. And like a normal router, VR will route IP traffic for the assigned customer.

In one embodiment, virtual private networks (VPNs) are logical abstractions. Each VPN represents the logical part of network visible or associated with a subscriber or ISP. A VPN primarily comprises of Virtual Routers (VR) and Virtual Interfaces (VI). In one such embodiment, service and performance characteristics of a VPN are captured in a Service Level Agreement (SLA) associated with the VPN.

In one embodiment, a virtual router can have multiple interfaces associated with it. These interfaces are termed Virtual Network Connections (VNC) or Virtual Interfaces (VI). A VNC or VI corresponds to only one side of a network connection. Virtual Interfaces represent either connection end points for a subscriber or a connection between two virtual routers. A VI mostly captures the network layer properties of an interface (one example of such a property of a VI is its IP address). A VI is mapped to either a Logical Interface or another VI. In one such embodiment, each VNC/VI is characterized by bandwidth, latency, services, and a set of prioritized queues.

VI Bandwidth is measured in bits per second (bps) and is the average number of bits per second transmitted over the VI during a sampling interval. Maximum absolute bandwidth may be set to enforce QOS commitments first then use additional capacity for overflow burst traffic. VI bandwidth can be a multiple of 56 Kbps or 64 Kbps.

VNC/VI Services include Application Firewall, Packet Filter, Encryption, NAT, Tunnel Termination, Tunnel origination and authentication (identification of the Virtual RADIUS Server), certificate authority, Accounting. Multiple VNC Services may be configured for the same VNC, directly from the VPN Manager.

VNC/VI Queues are used to prioritize different traffic types within a VNC. Each VNC Queue is characterized by bandwidth, latency, and traffic type. The sum of the bandwidth allocated on all VNC Queues for a particular VNC is equal to the VNC Bandwidth for that VNC. IP data traffic authorized for a particular VPN Queue is identified by using one or more of the following parameters: Source IP Network Address, Destination IP Network Address, IP Protocol Type, Port # and the Differentiated Services Byte.

A VI may be mapped into a Logical Interface (LI). A LI can be a PPP link, a PVC or SVC of FrameRelay or ATM, a Ethernet or an IP tunnel etc. A LI may be

mapped into a physical interface; it is not mandatory for every LI to be mapped to a physical interface.

In one embodiment, a subscriber's access to each VPN is captured using a Logical Interface (LI) terminating on an Access Blade port. One or more VRs are created and associated with a VPN to route subscribers IP traffic as per the subscriber requirements. In one embodiment, a subscriber's LI terminating on an Access Blade port is linked with a VR using a Virtual Interface (VI). One or more virtual network connections (VNCs) are associated with VIs to define subscriber's connectivity to the other end, which for example could be internet, intranet, subscribers regional office etc.

For instance, Internet Access VNCs may be used to connect a subscriber's VPN Router to the ISP's Internet backbone, thereby providing users at the subscriber premise with secure, dedicated Internet access.

VPN Intranet VNCs may be used to connect different subscriber locations securely across an IP backbone or the Internet.

VPN Extranet VNCs may be used to connect a subscriber's VPN Router to the VPN Router of another corporation or community of interest to allow access between the entities.

VPN RAS VNCs may be used to connect a subscriber's VPN router to an ISP's dial-up remote access server (such as Ascend MAX, Cisco 5800, etc.). The remote access servers (RAS) and/or the dial-in client initiates either a PPTP or L2TP tunnel to the Orion, with or without encryption.

The way that virtual routers map on processor elements can be understood in the context of Figure 4. Figure 4 shows a blade 28 having four processor elements (PEs) 30. Each processor element 30 includes a CPU and memory. In addition, each PE 30 is connected to the other PEs 30 and to the rest of switch 12 through connection fabric 32. In the embodiment shown in Figure 4, two virtual routers 34 are mapped on blade 28. It should be noted that, in one embodiment, more than one virtual router can be mapped to a single PE 30, and vice versa. Blade 28 can, therefore, be a shared resource among two or more subscribers.

The segmenting of switch 12 across a number of different subscribers is shown

in Figure 6.

In one embodiment, each VR is divided into a discrete object group, which has many functions and services associated with it; for instance routing, packet filtering and firewall each are individual objects. By being able to create a new object such as Network Address Translation (NAT), it is simply added as another part of the object group and activated. This action can be achieved without making any changes to the physical configuration of the Service Blade supporting the VR and thus maintaining network uptime.

The configuration of Virtual Routers (VR) and their corresponding Virtual Interfaces (VI) provides the desired constraints (or paths) for packet flows and the permissible transformations (AKA services) of the packet flows. NOS 20 still, however, has to provide mechanisms to efficiently direct packets from ingress port through protocol stack to Application service thread. The OS also has to provide mechanisms to efficiently direct packets from Application service thread through protocol stack to egress port.

In one embodiment, the configured VR/VI system defines a semi-static topology within the device boundary (which hereafter is referred to as the *Configured Topology*). Based on traffic patterns, the configured topology could induce a dynamic topology comprising of shortcuts (which hereafter is referred to as the *Flow Topology*).

In one embodiment, flows strictly follow the configured paths. This results in inefficient utilization of inter-connect bandwidth, memory and CPU time.

In another embodiment, NOS 20 supports a dynamic flow based topology as discussed below.

The nominal Protocol Stack configuration places topological constraints on packet flows. First, we will describe the *standard* protocol profile and related topology. Then we will covers the issues that arise when the protocol stack is distributed.

The following terms are used throughout:

Packet Flow All packets with a common set of attributes comprises a packet flow. Depending on the attributes, the flow may be L3 or L4 flow.

L3 Flow Only Layer 3 attributes are used to identify a packet flow. The Layer 3 attributes used are *<Destination IP Address, Source IP Address, IP Protocol>*.

L4 Flow When possible, both Layer 3 and Layer 4 attributes are used to identify a packet flow. Typically, multiple L4 flows comprises a single L3 flow. If the IP Protocol

is TCP or UDP the Layer 4 attributes used are *<Destination Port, Source Port>*.

Routing Table A common table of routing information that is used by Routing Protocols (e.g. BGP, OSPF, ISIS, RIP). The entries are usually tagged with origin, routing protocol attributes etc. This is also known as RIB (Routing Information Base).

Forwarding Table A table of best next-hop for prefixes, consulted in the packet forwarding function and derived from the Routing Table. This is also known as FIB (Forwarding Information Base).

Configured Topology Given a Subscriber VR and it's ISP VR, consider the associated graph. This graph's nodes are the PE's executing software/ services and the edges of the graph are the communication links between the software/ services. This graph is the configured topology.

Flow Topology Based on the actual dynamic flow of L3 and L4 packets a graph of software/ services and the communication links between the software/ services can be constructed. This graph is the flow topology. The flow topology is derived from and would be a sub-graph of the configured topology if shortcuts are not employed. When shortcuts are employed, there is no direct graph theoretic correlation between the configured and flow topologies.

Topological Constraints imposed by Protocol Stack

The Network layer is the loci of all packets received, sent or forwarded. Conceptually, the network layer (i.e. IP itself) can be thought of as being 2 sub-layers, namely the Forwarding and Local sub-layers. The functionality encompassed by the 2 IP sub-layers is shown in Figure 7. IP Interfaces are the link service access points (LSAP in OSI-speak). Traditional networking wisdom recommends:

For IP, a packet flow is identified by *<Source, Destination, Protocol>*. We will refer to these as *L3 flows*. These attributes of a packet are preserved across all packet fragments. Preserve the temporal order of L3 flows received on an IP Interface. This reduces the resequencing delay of stream based protocols (such as TCP). Many datagram based applications are intolerant of packet re-ordering (which is not unreasonable if the application is transported via a UDP tunnel without order preservation guarantees). The IP Packet ID, which is critical for packet fragmentation and reassembly, should be unique within a L3 flow originating from local applications. (This is actually required by RFC-791, Section 2.3, page 8).

Traditional implementations generally impose stronger restrictions than required, by:

Ensuring strict temporal order across all L3 flows received on an IP Interface.

Ensuring unique IP Packet ID across all L3 flows originating from local applications.

Packet forwarding needs to consult the *Forwarding Table* or *Forwarding Information Base (FIB)*. Most implementations maintain a single forwarding table. High performance implementations typically maintain multiple synchronized instances of the forwarding table.

We now look at what would happen in IPNOS 1.x for typical application flows. The inter-module transfers are shown in Figure 8. The processing sequence is

Packet is received from the subscriber line interface

Processed by IP Fwd module and passed to IP Local module

Demux'ed by IP Local and passed on to the Application

Application process data and sends transformed data

IP Local sends packet to subscriber IP Fwd module
Subscriber IP Fwd forwards it to ISP IP Fwd module
ISP IP Fwd module sends it to ISP line interface

Note that steps 2, 3, 4, 5, 6 require an inter-PE transfer. Each of these steps may represent a backplane crossing.

Figure 9 shows the contents that can be used to determine a L4 packet flow when the packet is unfragmented (determined by the condition ($MF == 0 \ \&\& \ Offset == 0$)). It also shows the contents that can be used to determine a L3 packet flow when the packet is fragmented.

When a packet gets fragmented, it should be noted that the ID field is preserved across all fragments of the packet. However, the receiver may receive the fragments in arbitrary order. Thus the receiver may only use <Source, Destination, Protocol> to determine the packet flow. The receiver can not create per packet flow state which uses the ID field to decide on the L4 flow for all fragments on the false assumption that the first fragment contains L4 information.

NOTE: All fragmented packets must be reassembled before the L4 flow can be inferred. This has implications for L4 oriented services like NAT. For NAT to work correctly, the gateway must reassemble every fragmented packet before it can perform its packet transformation. This is a tremendous burden and many NAT implementations (including IPSX) simply do not transform the fragments subsequent to the first.

There are several specific flows that we wish to have optimized. The forward flows are shown as originating at a Subscriber interface and ending at an ISP interface in the upper half of Figures 10-13. The reverse flow is shown in the lower half of each figure. Unoptimized flows are shown in black arrows and represent a fragment of the Configured Topology. The red arrows represent the optimized flows with shortcuts that we would like to generate and represent a fragment of the Flow Topology.

In one embodiment of NOS 20, objects communicate with each other in 1 of 2 ways – namely Object IOCTLs and Object Channels. The Object Channel is point-to-point (P-P), light weight inter-object communication transport. This maps well into the Configured Topology.

When shortcuts are created, an Object Channel termination will receive traffic from multiple Object Channel sources. That is, we need multi-point-to-point (MP-P) Object Channel in the Flow Topology. The PRD for MP-P describes the proposed implementation in detail.

The essential idea behind the MP-P Object Channel is that MP-P can be constructed in parallel to the P-P Object Channel. This permits the co-existence of the 2 communication models. This is shown in Figure 14.

Both ends of a P-P uses the initialization function `obj_init_channel()`. For MP-P, there will be 2 separate initialization functions. For MP-P transmit the initialization function is `obj_mpp_tx_init_channel()`. For MP-P receive the initialization function is `obj_mpp_init_channel()`. (Please refer to the PRD for actual function names). There is also an anonymous `mpp_send()` function that can be used (and which is bound to the channel end point by `obj_mpp_tx_init_channel()`) to send anonymously to a channel end point without using the Object Channel abstraction.

In Figure 14, Object Channel End Point (CEP) of both P-P and MP-P are shown. CEP 1

is initialized as both an MP-P receiver and as a P-P receiver/ sender. When CEP 1 was initialized as an MP-P receiver, it got assigned a globally unique address. CEP 2 is initialized solely as a P-P receiver/ sender. CEP 10 and 11 are initialized as MP-P sender with target address of CEP 1's MP-P address. Note that a MP-P sender *can not* be initialized as a P-P receiver/ sender. Additionally, modules that have CEP 1's MP-P address can send anonymously.

Another approach is to use a replicated FIB. The IP Forwarding agent uses the Forwarding Information Base (FIB) to make its forwarding decision. When the Virtual Router (VR) terminates links on multiple blades, the VR has to have multiple IP Forwarding agents. This requires the IP Forwarding agents to maintain a Replicated FIB. This feature has been referred to variously as *Multi-FIB*, *Distributed FIB* and *Split FIB* ... none of which reflects the actual operation.

The Replicated FIB need not be constrained to only IP Forwarding agents within a VR. A Replicated FIB may be used by other objects of a VR to make forwarding decisions to support shortcuts.

It should be noted that there is a host other data structures that also need to be replicated. These are Packet Filters (PF) (both dynamic and static), the Security Policy Database (SPD) and Security Association Database (SAD), and Network Address Translation (NAT) tables.

Another approach is to use a Distributed Protocol Stack. Figure 8 shows one possible functional decomposition of NOS 20 functionally distributed protocol stack. The functional decomposition simply cuts the protocol stack at convenient horizontal planes with no duplication of function.

To support a Flow Topology, a different functional decomposition of the protocol stack is required. This will duplicate functionality in multiple objects. Table 1 contrasts the assignment of functionality to objects under the two models of Distributed Protocol Stack.

It is a logical extension of the VR flow cache that is employed to speed up packet forwarding. It must be noted at the outset that this solution is significantly more complex than any previous proposals. Unfortunately, there does not seem to be any less complex solution available.

Objects in VR	IPNOS 1.1 Distributed Protocol Stack (Configured Topology)	Proposed Distributed Protocol Stack (Flow Topology)
IP Forwarding	IP input, IP output, IP Forwarding, IP Reassembly (ESP and AH only), IP Demux (ESP and AH only)	IP input, IP output, IP Forwarding, IP Reassembly (ESP and AH only), IP Demux (ESP and AH only)
IP Local	IP Reassembly, IP receive and send, IP Demux, TCP, UDP, GRE, Socket Lower, Routing protocols	IP input, IP output, IP Reassembly, IP receive and send, IP Demux, TCP, UDP, GRE, Socket Lower, Routing protocols
Crypto	Crypto services	IP input, IP output, IP Forwarding, IP Reassembly (ESP and AH only), IP Demux (ESP and AH only), Crypto services
Socket Upper	Socket Upper	IP input, IP output, IP Reassembly, IP receive and send, IP Demux, TCP, UDP, GRE, Socket Lower, Socket Upper
Application	Application	Application

Table 1: Functional Decomposition of Distributed Protocol Stacks

In the Flow based model, there are multiple objects that implement the IP functions. This requires that the FIB be replicated not only at the IP Forwarding object, but also in a host of other objects that implement the IP Forwarding layer functionality.

Distributed Adaptive Flow Cache

When a packet is received, it is looked up in the local L3 and if appropriate L4 flow cache. If either a L3 or L4 flow cache entry is found, it is forwarded unmarked. If no flow cache is found, then the packet is marked as "*L3/L4 cache information is requested*", and forwarded along the Configured Topology. Agents send shortcut information in the opposite direction to the source of the packets. This shortcut information gets installed and induces the Flow Topology.

The flow cache is periodically flushed to avoid stale references. Additionally, a refresh message is sent periodically to maintain established flow cache state.

Operational Example

To provide an operational overview of how Distributed Protocol Stack with Flow

Caching works, the sequence of operations, the flow state that is created, and the shortcuts that are established are shown in a sequence of figures. The specific flow that is used to explain the operational behavior is the reverse flow of **Error! Reference source not found.** The functional blocks that participate in the reverse flow is reproduced in Figure 1.

Figure 1 shows the packet flow for the 1st uncached packet. Figure 2, Figure 3 and Figure 4 indicate the cache state advertisements generated on the 1st packet.

Figure 5 shows the packet flow for the 2nd packet as a result of the cache advertisements induced by the 1st packet. Figure 6 indicates the cache advertisement generated on the 2nd packet.

Figure 7 shows the packet flow for the 3rd packet as a result of the cache advertisements induced by the 1st and 2nd packets.

Note that each object in the processing stages informs the earlier object of the short cut. Thus if the processing was A → B → C, B informs A about the shortcut to C (i.e. A → C).

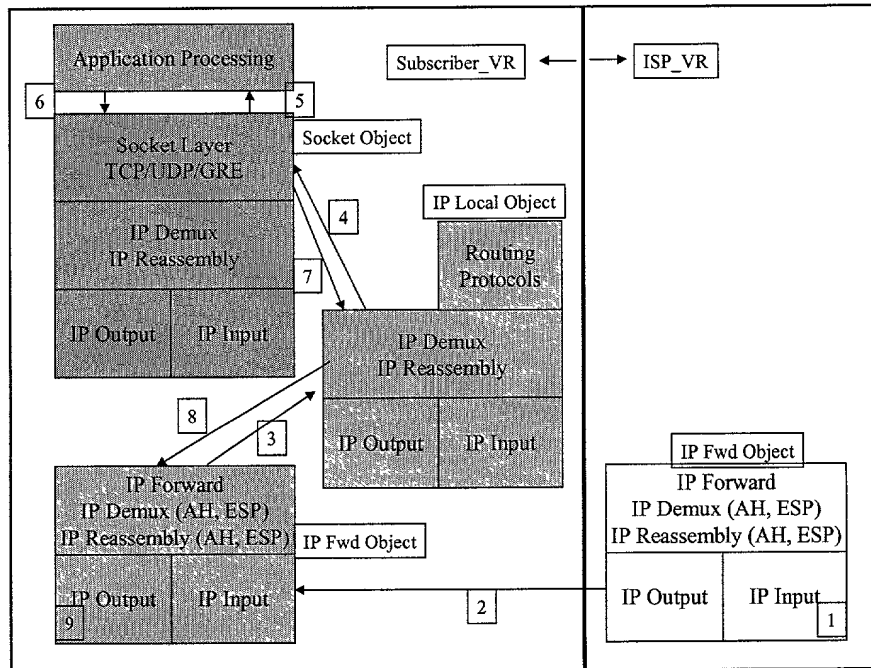


Figure 1: Uncached reverse socket flow for 1st packet

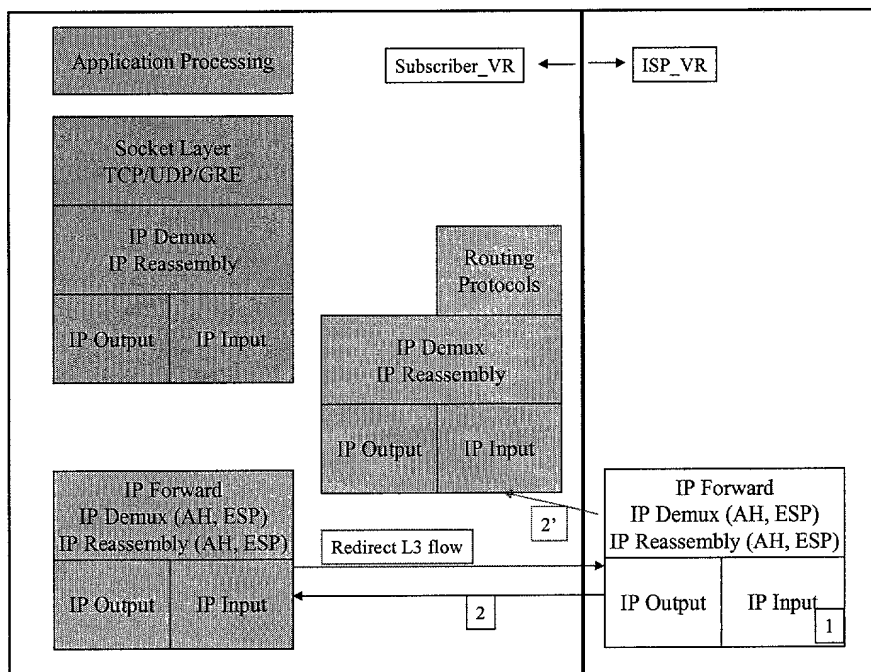


Figure 2: Input L3 cache state installed in ISP VR's IP Fwd Object on 1st packet

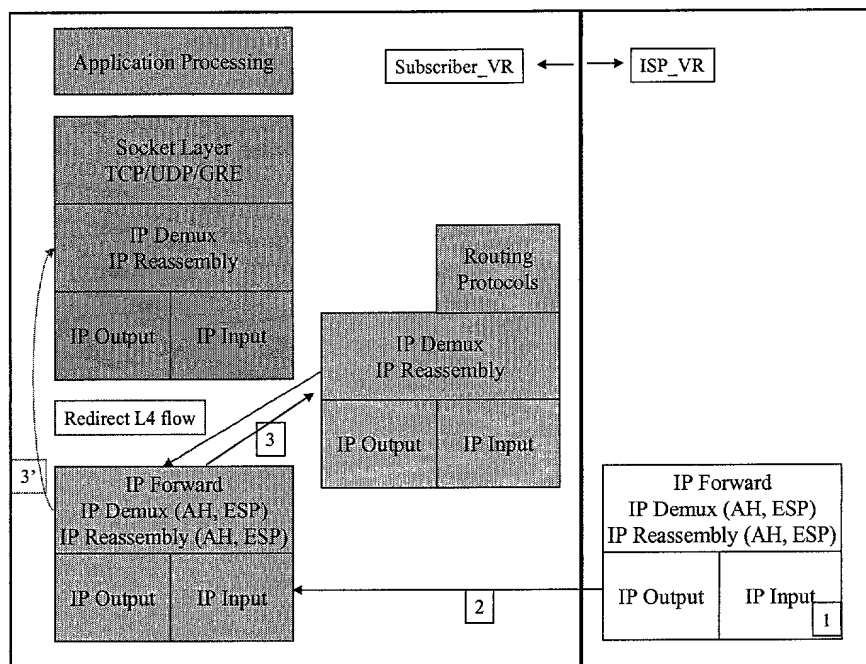


Figure 3: Input L4 cache state installed in Subscriber VR's IP Fwd Object on 1st packet

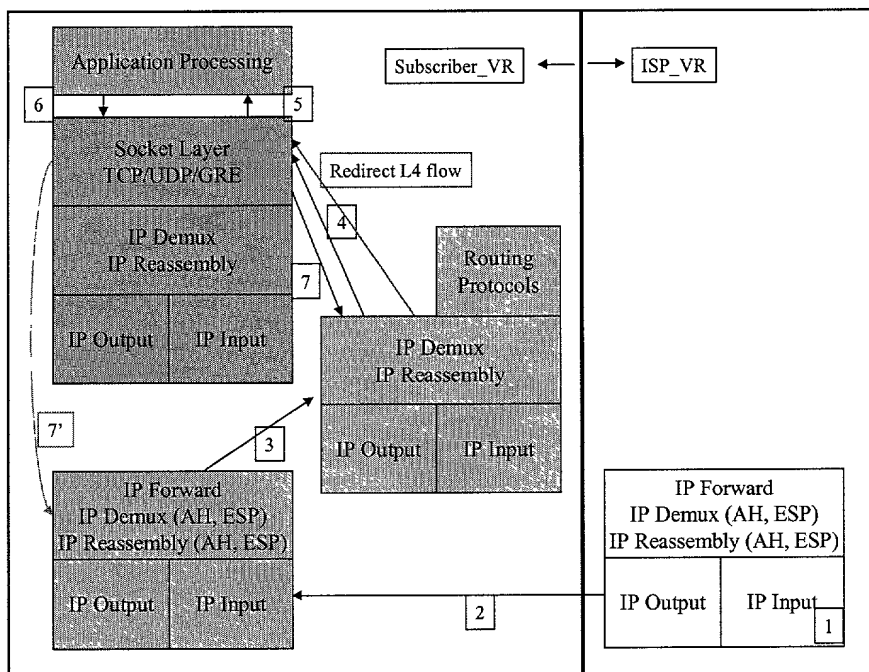


Figure 4: Output L4 cache state installed in Subscriber VR's Socket Object on 1st packet

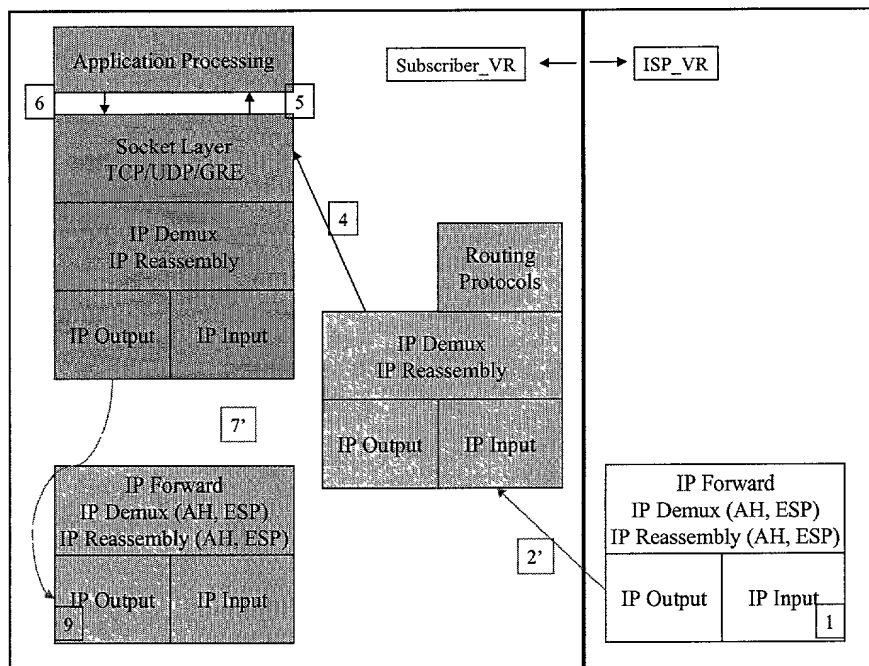


Figure 5: Packet flow for 2nd packet

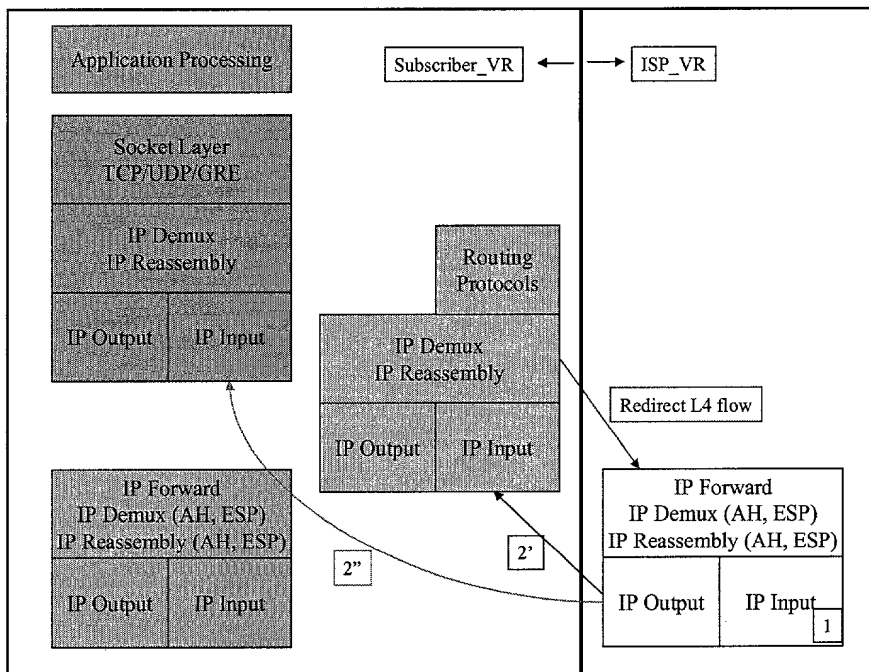


Figure 6: Input L4 cache state installed in ISP VR's IP Fwd Object on 2nd packet

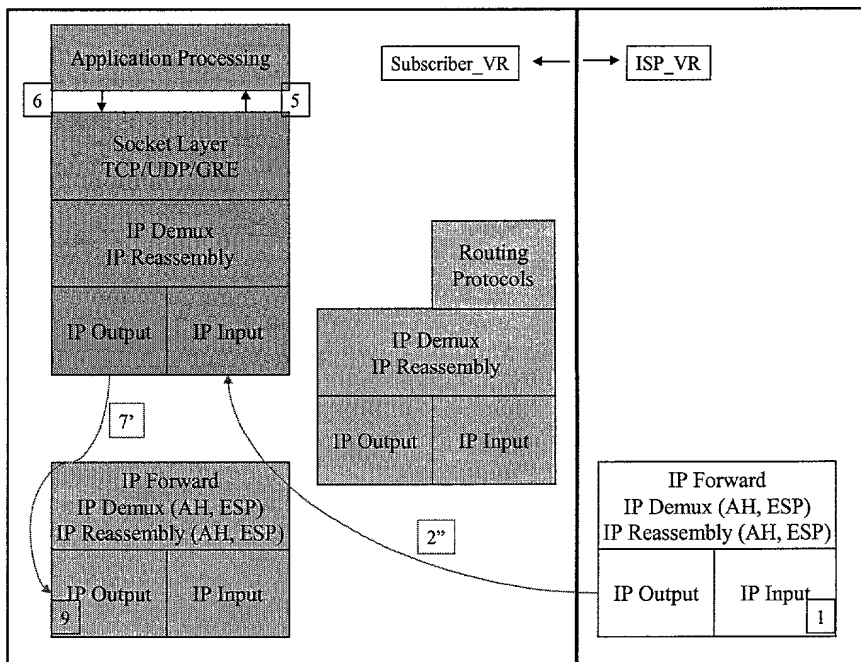


Figure 7: Packet flow for 3rd packet

1.1.1.1.Open Issues

This section covers some issues that do not have adequate definition. The issues are presented in no particular order.

- 1) The Object Identifiers (OID) used by the Object Manager (OM) is a multi-octet value that is rather large. For flow cache requests to be indicated and responded to, objects may need to be assigned a Run Time OID (RTOID) as well. This RTOID must be a 32-bit globally unique identifier.
- 2) It is worth noting that:
 - a) The distribution of information crosses VR boundaries.
 - b) The leaking of information across VR boundaries is constrained by the interconnection between VRs. Each object determines the extent to which such information is leaked and should ensure that it is appropriate.
 - c) Flow state is soft but is cached for a certain period.
 - d) Network topology changes that are more rapid than the flow cache retention period will lead to forwarding loops.
- 3) What is the effect of packets arriving faster than the rate at which cache information is flowing in the reverse direction? Can the device melt-down due to increased cache advertisements?
- 4) Do we really need functions like IP Input, IP Output etc. in objects other than IP Forward? They have been thrown in at this time ... but this needs to be evaluated carefully. In general, the functional decomposition needs further study.
- 5) Objects have multiple entry points based on the type of shortcut that is employed. E.g. we need 1 entry point from below for no shortcuts, 1 entry point from below for 1 shortcut, 1 entry point from below for 2 shortcuts etc. How do we
 - e) Indicate the entry point when sending cache advertisements?
 - f) How do we indicate the entry point when sending packets?
- 6) The protocol, flow state manager etc. need precise definition. It will be addressed once there is agreement that this is a valid approach.

1.2. Applicability of Solutions to Issues

Table 2 details the applicability of each solution to the issues. Each issue is broken down into a Forward and a Reverse direction for clarity. This table is presented without an explanation for each entry (and is also expected to be well debated!).

Solutions	Encryption Flows		Socket Flows		Distributed Forwarding		Tunneled Flows	
	F	R	F	R	F	R	F	R
Multi point to point	√	X	X	X	√	X	X	X
Replicated FIB	X	X	X	X	X	X	X	X
Distributed Protocol Stack	X	X	X	X	X	X	X	X
Multi point to point + Replicated FIB	√	X	√	X	√	√	X	X
Distributed Protocol Stack + Multi point to point	√	√	√	√	X	X	√	√
Distributed Protocol Stack + Replicated FIB	X	X	X	X	X	X	X	X
Distributed Protocol Stack + Multi point to point + Replicated FIB	√	√	√	√	√	√	√	√

Table 2: Applicability of solutions to optimizing issues

In one embodiment, Object Manager 24 consists of three layers as shown on Fig. 4. The upper layer titled *OM Controller and Database* (OMCD) 40 is concerned with managing the VPN and VR configuration. This is the agent that deals with the configuration manager directly. Middle layer 42 entitled *OM Object Routing and Interface Global* is concerned with managing global (across the switch system) object groups and object configurations. Lower layer 44 entitled *OM Object Routing and Interface* (OMORI) is concerned with managing local objects and groups as well as routing control information between address spaces based on the location of objects, and interfacing with the object via method invocation.

In one embodiment, the IPSX object database consists of two types of databases: Global (managed on Master Control Blade by OMORIG) and distributed local databases (managed by OMORI agents on every PE present in the system). In one such embodiment, the global database is a superset of the extracts from local databases.

Objects represent a basic unit of management for purposes of fault tolerance, computational load balancing, etc. One or more adjacent protocol modules can be

placed into a single object. It is also possible that a module is split across two objects.

In one embodiment, a configuration management system connected to OMCD 40 is used to configure the VPNs and VRs as shown in Appendix A and is described in "System and Method for Configuration Management", filed herewith, the description of which is incorporated herein by reference.

In one embodiment, instead of deploying CPE firewalls remotely at the enterprise subscriber site, System 10 provides Network Associates, Incorporated's (NAI's) Gauntlet Application Proxy Firewalls from the SP POP where they can be centrally managed and supported. NOS 20 is the foundation that enables third-party applications like Gauntlet to be delivered on switch 12.

The application proxy firewall model offers far superior security control because it provides full application-level awareness of attempted connections by examining everything at the highest layer of the protocol stack.

Once NAI's Gauntlet Application Proxy Firewall is configured, it offers the most robust form of access control on the market. Gauntlet lets access control policies to be established at Layer 7 and thus enables specific application-based rules to be made. For instance, users coming into the network from the Internet are allowed to do an FTP "get," but not a "put." Additionally, because the Gauntlet firewall is a proxy solution that acts as the intermediary, preventing the end-stations from communicating directly with the firewall. As a result, this solution offers significantly better protection against OS directed denial of service attacks. Another important aspect of the Application Proxy Firewall is the incremental services provided by switch 12 such as Intrusion Detection and Anti-virus scanning, which are not available in firewall application models such as stateful inspection.

Conclusion

In one embodiment, the service provider's security staff consults with the customer in order to understand the corporate network infrastructure and to develop appropriate security policies (Note: this is a similar process to the CPE model). Once this has been accomplished, the NOC security staff remotely accesses the IP Service

Processing Switch (using the Service Management System) at the regional POP serving the enterprise customer, and the firewall service is provisioned and configured remotely. Such a model enables the service provider to leverage the enterprise's existing services infrastructure (leased lines and Frame Relay PVCs) to deliver new, value-added services without the requirement of a truck roll. All firewall and VPN functionality resides on the IP Service Processing Switch at the POP, thus freeing the service provider from onsite systems integration and configuration and effectively hiding the technology from the enterprise customer. Firewall inspection and access control functions, as well as VPN tunneling and encryption, take place at the IP Service Processing Switch and across the WAN, while the enterprise's secure leased line or Frame Relay PVC access link remains in place. The customer interface is between its router and the IP Service Processing Switch (acting as an access router), just as it was prior to the rollout of the managed firewall service. Additionally, the customer has visibility into and control over its segment of the network via the CNM that typically resides at the headquarters site.

The network-based firewall model also enables service providers to quickly and cost-effectively roll out managed firewall solutions at all enterprise customer sites. As a result, secure Internet access can be provided to every site, eliminating the performance and complexity issues associated with backhauling Internet traffic across the WAN to and from a centralized secure access point. As the IP Service Delivery Platform is designed to enable value-added public network services, it is a carrier-grade system that is more robust and higher-capacity than traditional access routers, and an order of magnitude more scaleable and manageable than CPE-based systems. The platform's Service Management System enables managed firewall services, as well as a host of other managed network services, to be provisioned, configured, and managed with point-and-click simplicity, minimizing the need for expensive, highly skilled security professionals and significantly cutting service rollout lead-times. The Service Management System is capable of supporting a fleet of IP Service Processing Switches and tens of thousands of enterprise networks, and interfaces to the platform at the POP from the NOC via IP address. Support for incremental additional platforms and customers is added via modular software add-ons. Services can be provisioned via the

SMS system's simple point and click menus, as well as requested directly by the customer via the CNM system. Deployment of a robust IP Service Delivery Platform in the carrier network enables service providers to rapidly turn-up high value, managed network-based services at a fraction of the capital and operational costs of CPE-based solutions. This enables service providers to gain a least-cost service delivery and support structure. Additionally, it enables them to gain higher margins and more market share than competitors utilizing traditional service delivery mechanisms — even while offering managed firewall services at a lower customer price point.

As enterprise customers gain confidence in the WAN providers' ability to deliver managed firewall services, a more scaleable and cost-effective service delivery model must be employed. Moving the intelligence of the service off of the customer premises and into the WAN is an effective strategy to accomplish this. Managed, network-based firewall services provide the same feature/functionality of a CPE-based service while greatly reducing capital and operational costs, as well as complexity.

The managed, network-based firewall service model enables WAN service providers to minimize service creation and delivery costs. This model virtually eliminates the need for onsite installation, configuration, and troubleshooting truck rolls, drastically reducing operational costs. This lower cost structure creates opportunities to increase revenues and/or gain market share by value-pricing the service. Services can be rapidly provisioned via a centralized services management system, shortening delivery cycles and enabling service providers to begin billing immediately. Additional services can be rapidly crafted and deployed via the same efficient delivery mechanism.

The network-based service model is a rapid and cost-effective way for service providers to deploy high-value managed firewall solutions. This model requires a comprehensive service delivery platform consisting of robust network hardware, an intelligent and scaleable services management system, and a feature-rich Customer Network Management (CNM) tool to mitigate customers' fears of losing control of network security.

In the above discussion and in the attached appendices, the term "computer" is defined to include any digital or analog data processing unit. Examples include any

personal computer, workstation, set top box, mainframe, server, supercomputer, laptop or personal digital assistant capable of embodying the inventions described herein.

Examples of articles comprising computer readable media are floppy disks, hard drives, CD-ROM or DVD media or any other read-write or read-only memory device.

Although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiment shown. This application is intended to cover any adaptations or variations of the present invention. Therefore, it is intended that this invention be limited only by the claims and the equivalents thereof.

What is claimed is:

1. A virtual routing system, comprising:
a virtual router operating across one or more processing elements.

SYSTEM AND METHOD FOR ALLOCATING NETWORKING RESOURCES

Abstract

A virtual routing system, including a virtual router operating across one or more processing elements.

"Express Mail" mailing label number: EL618477168US

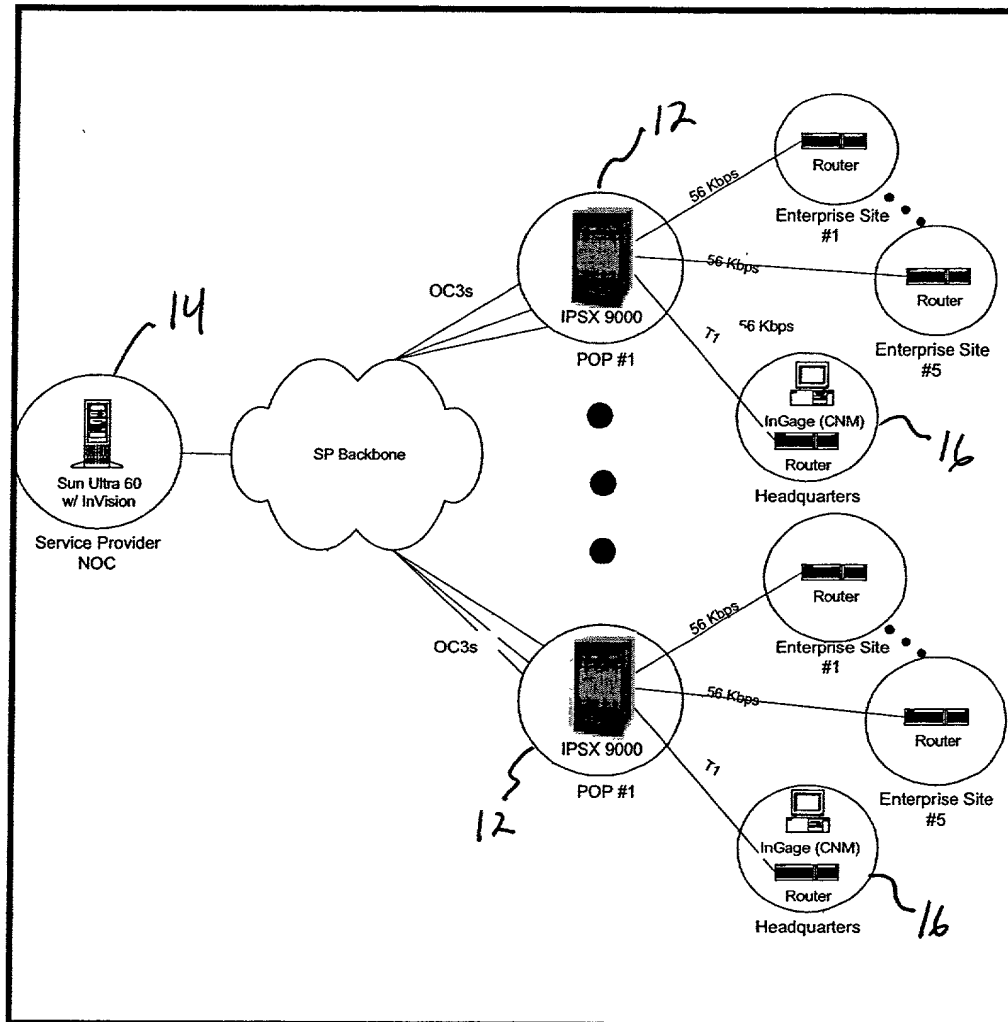
Date of Deposit: September 13, 2000

This paper or fee is being deposited on the date indicated above with the United States Postal Service pursuant to 37 CFR 1.10, and is addressed to the Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

10



Diagram 5 — Managed Firewall Service with CoSine's Network-based Solution



POP Infrastructure

The POP access infrastructure in the network-based managed firewall service model is based on the CoSine Communications IPSX 9000 Service Processing Switch. The base configuration for the switch includes:

- 26-slot chassis
- Redundant power supply
- IPNOS Base Software
- Ring Bridge & Ring Bridge Pass-Thru (to complete midplane)
- Control Blade (for communications with InVision Services Management System)
- Dual-port Channelized DS3 Access Blade
- Dual-port Unchannelized DS3 Access Blades
- Processor Blade
- OC-3c POS Trunk Blade

The following tables analyze the cost structure of all of the above models and projects these costs out over 5 years:

IPNOS Overview




Fig. 3

2. OM Design

2.1. Overview

The Object Manager consists of 3 layers as shown on Figure 1. The upper layer titled *OM Configuration Database (OMCD)* is concerned with managing the VPN and VR configuration. This is the agent that deals with CM directly. Middle layer titled is *OM Object Routing and Interface Global* concerned with managing global (across the IPSX system) object groups and object configurations. The lower layer titled *OM Object Routing and Interface (OMORI)* is concerned with managing local objects and groups as well as routing control information between address spaces based on the location of objects, and interfacing with the object via method invocation.

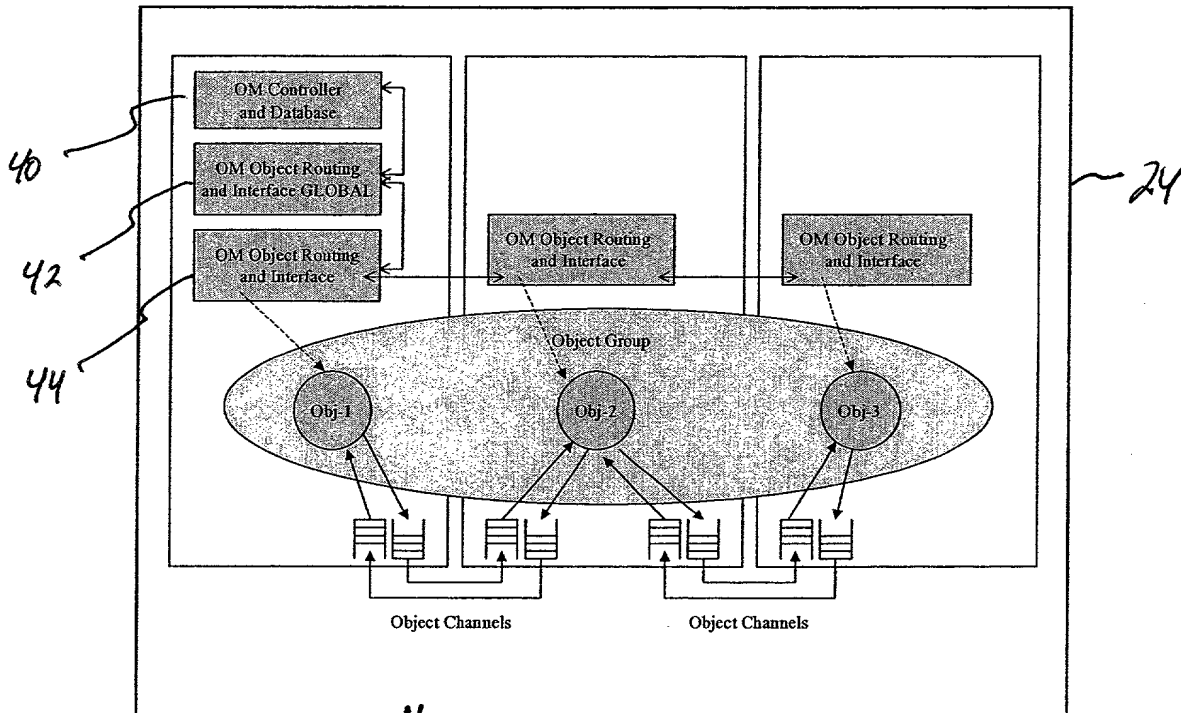


Figure 1. Object Manager Layers

2.1.1. Database distribution

IPSX object database consists of two types of database: Global, managed on Master Control Blade by OMORIG and distributed local databases, managed by OMORI agents on every PE present in the system. Global database is a superset of the extracts from local databases.

2.2. Object

Objects represent a basic unit of management for purposes of fault tolerance, computational load balancing etc. One or more adjacent protocol modules can be placed into a single object. It is also possible that a module is split across two objects.




Fig. 5

Innovative Method for Segmenting and Layering Services




Fig. 6

- Each subscriber has a set of partitioned Virtual Routers (VRs)
- Each VR is the equivalent of an independent hardware router
- VR as an object group enables customized services per subscriber
- InVision allows ease of service provisioning and maintenance of services across all IPSX units in a SP network
- IP Network Operating System's (IPNOS) open Application Program Interface (API) enables new services to be continually added to the platform

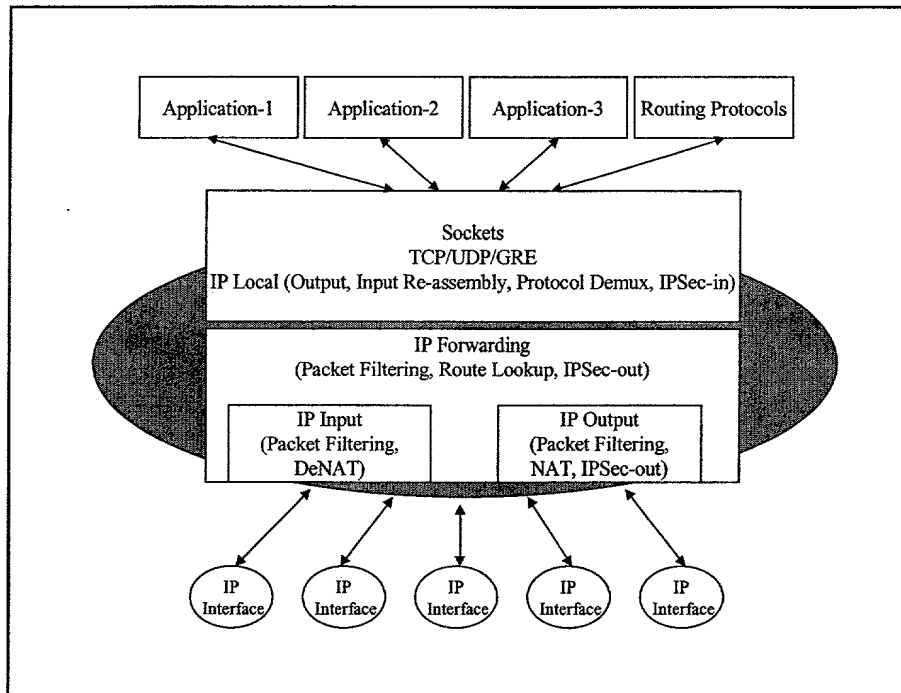


Figure 7: Standard Protocol Stack Profile

We now look at what would happen in IPNOS 1.x for typical application flows. The inter-module transfers are shown in Figure 8. The processing sequence is

- 1) Packet is received from the subscriber line interface
- 2) Processed by IP Fwd module and passed to IP Local module
- 3) Demux'ed by IP Local and passed on to the Application
- 4) Application process data and sends transformed data
- 5) IP Local sends packet to subscriber IP Fwd module
- 6) Subscriber IP Fwd forwards it to ISP IP Fwd module
- 7) ISP IP Fwd module sends it to ISP line interface

Note that steps 2, 3, 4, 5, 6 require an inter-PE transfer. Each of these steps may represent a backplane crossing.

CoSine Communications Inc.	Architecture Requirements Document	
Company Confidential. For Internal Circulation only	IPSX-P~1	4




Figure 3: Firewall flow processing

2.3 Differentiating between L3 and L4 flows

Figure 3 shows the contents that can be used to determine a L4 packet flow when the packet is unfragmented (determined by the condition $(MF == 0 \ \&\& \ Offset == 0)$). It also shows the contents that can be used to determine a L3 packet flow when the packet is fragmented.

When a packet gets fragmented, it should be noted that the ID field is preserved across all fragments of the packet. However, the receiver may receive the fragments in arbitrary order. Thus the receiver may only use $\langle Source, Destination, Protocol \rangle$ to determine the packet flow. The receiver can not create per packet flow state which uses the ID field to decide on the L4 flow for all fragments on the false assumption that the first fragment contains L4 information.

NOTE: All fragmented packets must be reassembled before the L4 flow can be inferred. This has implications for L4 oriented services like NAT. For NAT to work correctly, the gateway must reassemble every fragmented packet before it can perform its packet transformation. This is a tremendous burden and many NAT implementations (including IPSX) simply do not transform the fragments subsequent to the first.

CoSine Communications Inc.	Architecture Requirements Document	
Company Confidential. For Internal Circulation only	IPSX-P~1	5

Original Unfragmented IP Packet (MF= 0 && Off= 0)

Src, Dest, Proto MF=0, Off=0 ID=10AF, len=1064	Src Port Dest Port	Data[0-1023]
--	-----------------------	--------------

Fragmented IP Packet (each fragment is 534 bytes long)

Src, Dest, Proto MF=1, Off=0 ID=10AF, len=512	Src Port Dest Port	Data[0-471]	IP Fragment 1
Src, Dest, Proto MF=1, Off=492 ID=10AF, len=512		Data[472-963]	IP Fragment 2
Src, Dest, Proto MF=1, Off=984 ID=10AF, len=80		Data[964-1023]	IP Fragment 3

Figure 9: Packet Fragmentation and Header Content

3 Issues and Solutions

3.1 Issues to be addressed

In this section we identify several specific flows that we wish to have optimized. The forward flows are shown as originating at a Subscriber interface and ending at an ISP interface in the upper half of Figure 4, Figure 5, Figure 6 and Figure 7. The reverse flow is shown in the lower half of each figure.

The un-optimized flows are shown in black arrows and represent a fragment of the Configured Topology. The red arrows represent the optimized flows with shortcuts that we would like to generate and represent a fragment of the Flow Topology.

CoSine Communications Inc.	Architecture Requirements Document	
Company Confidential. For Internal Circulation only	IPSX-P~1	6




Figure 4: Encryption flows

10




Figure 5: Socket flows

11

CoSine Communications Inc.	Architecture Requirements Document	
Company Confidential. For Internal Circulation only	IPSX-P~1	7




Figure 6: Distributed Forwarding

12




Figure 7: Tunnelled flows

13

CoSine Communications Inc.	Architecture Requirements Document	
Company Confidential. For Internal Circulation only	IPSX-P~1	8

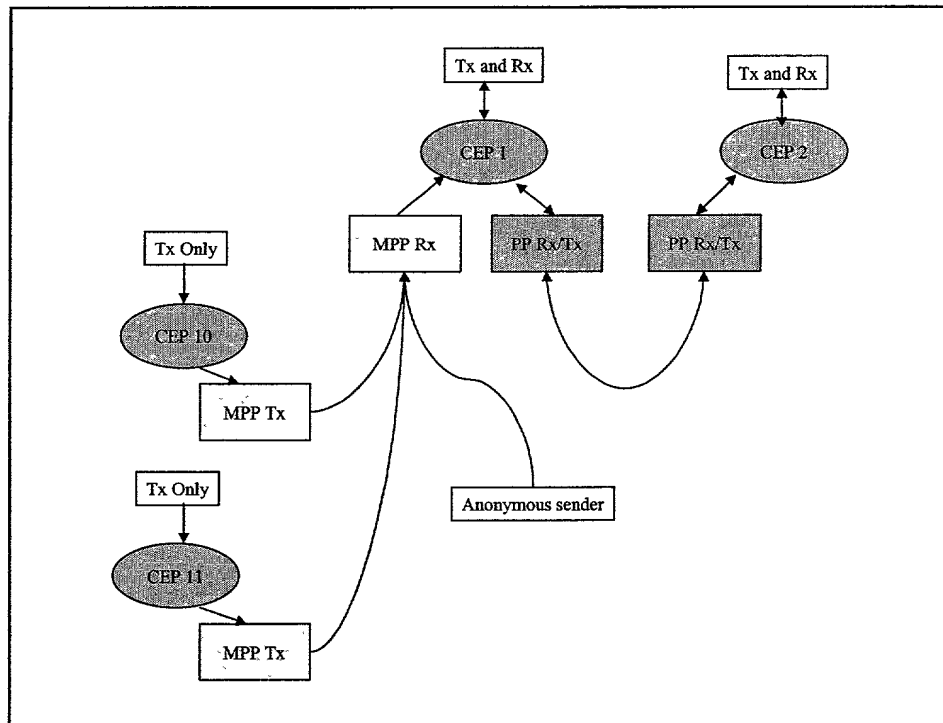


Figure 8: MP-P Operation

3.2.2 Replicated FIB

The IP Forwarding agent uses the Forwarding Information Base (FIB) to make its forwarding decision. When the Virtual Router (VR) terminates links on multiple blades, the VR has to have multiple IP Forwarding agents. This requires the IP Forwarding agents to maintain a Replicated FIB. This feature has been referred to variously as *Multi-FIB*, *Distributed FIB* and *Split FIB* ... none of which reflects the actual operation.

The Replicated FIB need not be constrained to only IP Forwarding agents within a VR. A Replicated FIB may be used by other objects of a VR to make forwarding decisions to support shortcuts.

It should be noted that there is a host other data structures that also need to be replicated. These are


- 1) Packet Filters (PF), both dynamic and static.
- 2) Security Policy Database (SPD) and Security Association Database (SAD).
- 3) Network Address Translation (NAT) tables.

3.2.3 Distributed Protocol Stack (Flow Based)

IPNOS 1.1 supports a functionally distributed protocol stack, which strictly implements the Configured Topology. Figure 2 shows the planned functional decomposition of the

CoSine Communications Inc.	Architecture Requirements Document	
Company Confidential. For Internal Circulation only	IPSX-P~1	10


IPSX Virtual Router Objects




CONFIDENTIAL INFORMATION

CONFIDENTIAL


IPSX Virtual Router Object Placement




IPSX Virtual Router Instances



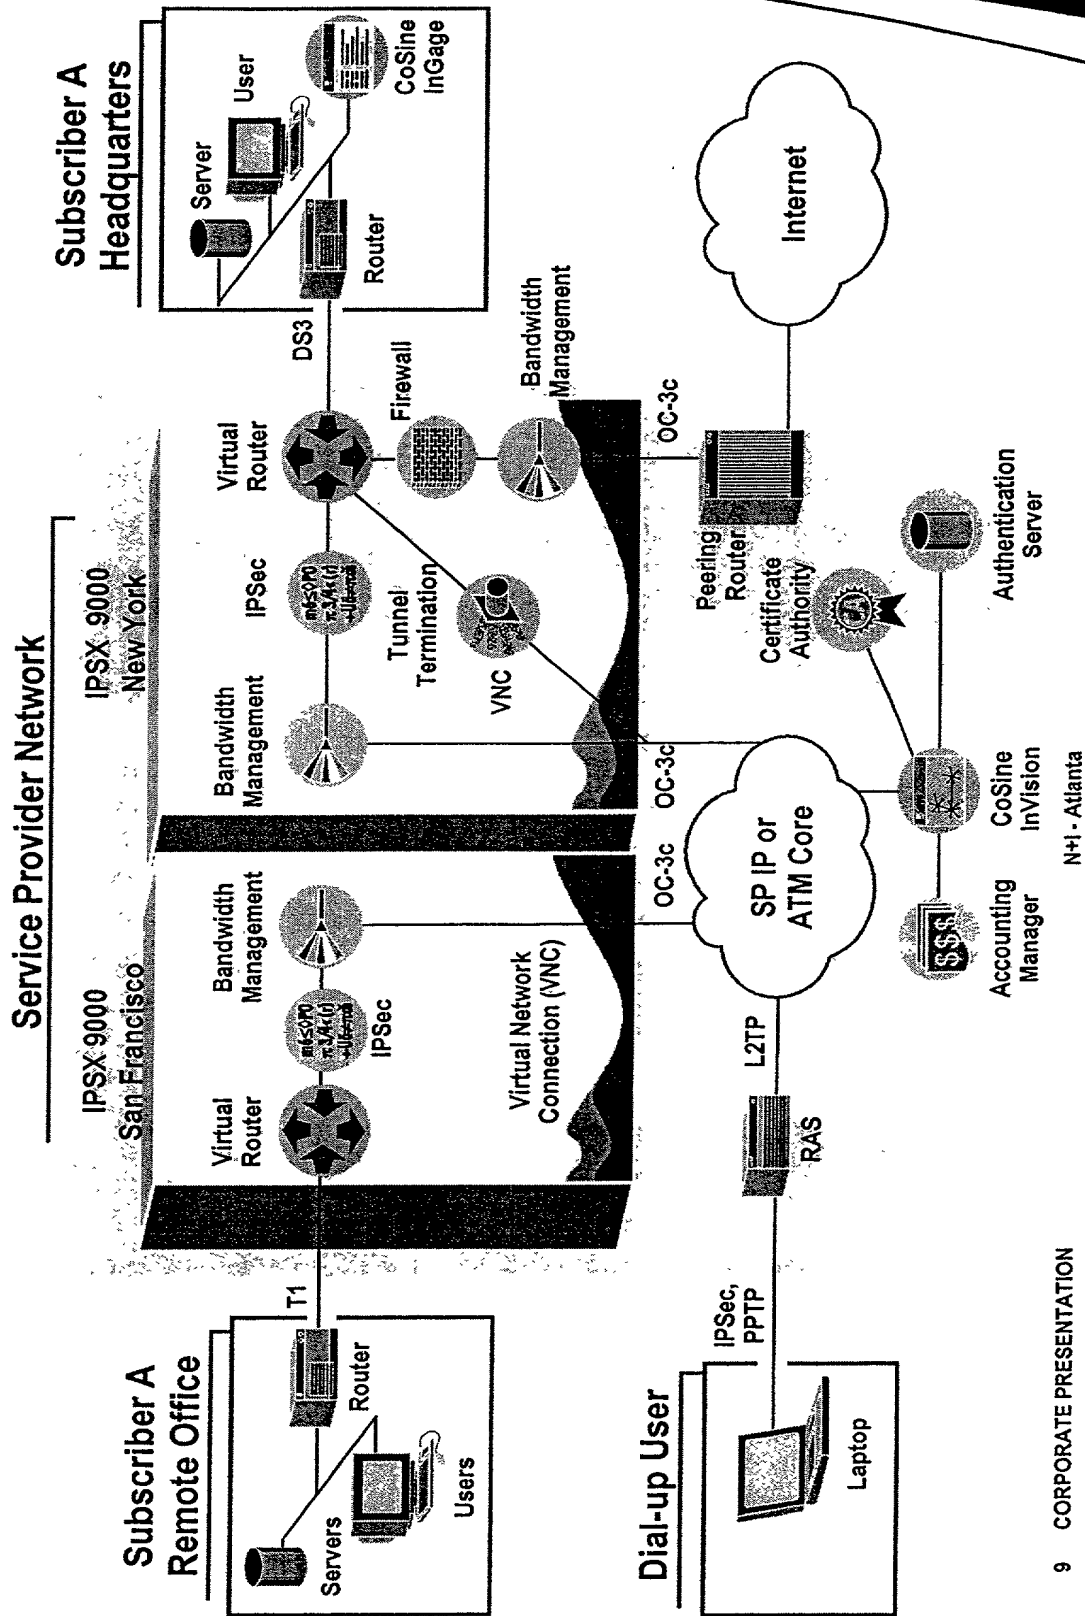
Managed, Service Provider-based VPNs: Intranet Connectivity and Secure Internet Access




Service Provider Adds Managed Dial-Up Services to VPN



SP Adds Managed Bandwidth Management Services to VPN



Dynamic SLAs: CoSine Virtual Network Connection (VNC) Detail



- Bandwidth Management service enables establishment of up to 32 rate shaped queues between virtual routers or on any given interface.
- Each VNCQ may have constituents assigned by user and application, for example:

VNCQ-1: All engineering FTP traffic – 64 Kbps

VNCQ-2: All executive officer traffic – 64 Kbps

VNCQ-n: All http traffic – 12 Kbps




Fig.

SCHWEGMAN ■ LUNDBERG ■ WOESSNER ■ KLUTH

United States Patent Application

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor I hereby declare that: my residence, post office address and citizenship are as stated below next to my name; that

I verily believe I am the original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled: **SYSTEM AND METHOD FOR ALLOCATING NETWORKING RESOURCES.**

The specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 C.F.R. § 1.56 (attached hereto). I also acknowledge my duty to disclose all information known to be material to patentability which became available between a filing date of a prior application and the national or PCT international filing date in the event this is a Continuation-In-Part application in accordance with 37 C.F.R. § 1.63(c).

I hereby claim foreign priority benefits under 35 U.S.C. § 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on the basis of which priority is claimed:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 120 or 365(c) of any United States and PCT international application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. § 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. § 1.56(a) which became available between the filing date of the prior application and the national or PCT international filing date of this application:

No such claim for priority is being made at this time.

I hereby appoint the following attorney(s) and/or patent agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith:

Anglin, J. Michael	Reg. No. 24,916	Huebsch, Joseph C.	Reg. No. 42,673	Nelson, Albin J.	Reg. No. 28,650
Bianchi, Timothy E.	Reg. No. 39,610	Jurkovich, Patti J.	Reg. No. 44,813	Nielsen, Walter W.	Reg. No. 25,539
Billion, Richard E.	Reg. No. 32,836	Kalis, Janal M.	Reg. No. 37,650	Oh, Allen J.	Reg. No. 42,047
Black, David W.	Reg. No. 42,331	Kaufmann, John D.	Reg. No. 24,017	Padys, Danny J.	Reg. No. 35,635
Brennan, Leoniede M.	Reg. No. 35,832	Klima-Silberg, Catherine I.	Reg. No. 40,052	Parker, J. Kevin	Reg. No. 33,024
Brennan, Thomas F.	Reg. No. 35,075	Kluth, Daniel J.	Reg. No. 32,146	Perdok, Monique M.	Reg. No. 42,989
Brooks, Edward J., III	Reg. No. 40,925	Lacy, Rodney L.	Reg. No. 41,136	Prout, William F.	Reg. No. 33,995
Chu, Dinh C.P.	Reg. No. 41,676	Lemaire, Charles A.	Reg. No. 36,198	Schumm, Sherry W.	Reg. No. 39,422
Clark, Barbara J.	Reg. No. 38,107	LeMoine, Dana B.	Reg. No. 40,062	Schwegman, Micheal L.	Reg. No. 25,816
Clise, Timothy B.	Reg. No. 40,957	Lundberg, Steven W.	Reg. No. 30,568	Scott, John C.	Reg. No. 38,613
Dahl, John M.	Reg. No. 44,639	Maeyaert, Paul L.	Reg. No. 40,076	Smith, Michael G.	Reg. No. 45,368
Drake, Eduardo E.	Reg. No. 40,594	Maki, Peter C.	Reg. No. 42,832	Speier, Gary J.	Reg. No. 45,458
Embretson, Janet E.	Reg. No. 39,665	Malen, Peter L.	Reg. No. 44,894	Steffey, Charles E.	Reg. No. 25,179
Fordenbacher, Paul J.	Reg. No. 42,546	Mates, Robert E.	Reg. No. 35,271	Terry, Kathleen R.	Reg. No. 31,884
Forrest, Bradley A.	Reg. No. 30,837	McCrackin, Ann M.	Reg. No. 42,858	Tong, Viet V.	Reg. No. 45,416
Gamon, Owen J.	Reg. No. 36,143	Moore, Charles L., Jr.	Reg. No. 33,742	Viksnins, Ann S.	Reg. No. 37,748
Harris, Robert J.	Reg. No. 37,346	Nama, Kash	Reg. No. 44,255	Woessner, Warren D.	Reg. No. 30,440

I hereby authorize them to act and rely on instructions from and communicate directly with the person/assignee/attorney/firm/organization/who/which first sends/sent this case to them and by whom/which I hereby declare that I have consented after full disclosure to be represented unless/until I instruct Schwegman, Lundberg, Woessner & Kluth, P.A. to the contrary.

Please direct all correspondence in this case to **Schwegman, Lundberg, Woessner & Kluth, P.A.** at the address indicated below:
P.O. Box 2938, Minneapolis, MN 55402
Telephone No. (612)373-6900

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of joint inventor number 1 : **Abraham R. Matthews**
Citizenship: **India** Residence: **San Jose, CA**
Post Office Address: **933 Willowleaf Drive**
Apt. 605
San Jose, CA 95128

Signature: _____ Date: _____
Abraham R. Matthews

Full Name of joint inventor number 2 : **Anna Berenberg**
Citizenship: **United States of America** Residence: **Redwood City, CA**
Post Office Address: **3200 Bridge Parkway**
Redwood City, CA 94065

Signature: _____ Date: _____
Anna Berenberg

☒ Additional inventors are being named on separately numbered sheets, attached hereto.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of joint inventor number 3 : **Kevin Lin**
Citizenship: **United States of America** Residence: **Redwood City, CA**
Post Office Address: **3200 Bridge Parkway**
Redwood City, CA 94065

Signature: _____ Date: _____
Kevin Lin

Full Name of inventor:
Citizenship: Residence:
Post Office Address:

Signature: _____ Date: _____

Full Name of inventor:
Citizenship: Residence:
Post Office Address:

Signature: _____ Date: _____

Full Name of inventor:
Citizenship: Residence:
Post Office Address:

Signature: _____ Date: _____

§ 1.56 Duty to disclose information material to patentability.

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is canceled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is canceled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) the closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

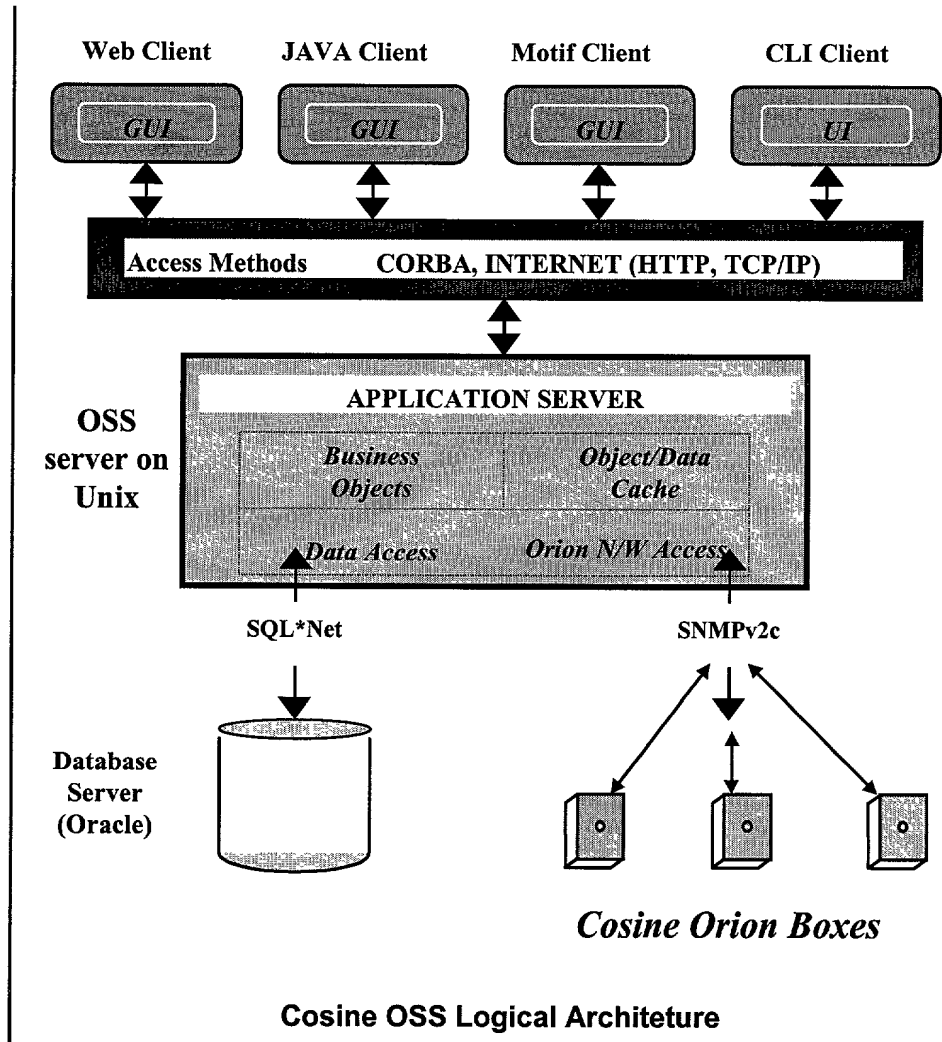
(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.

Appendix A

4. Other Requirements

4.1 Logical Architecture

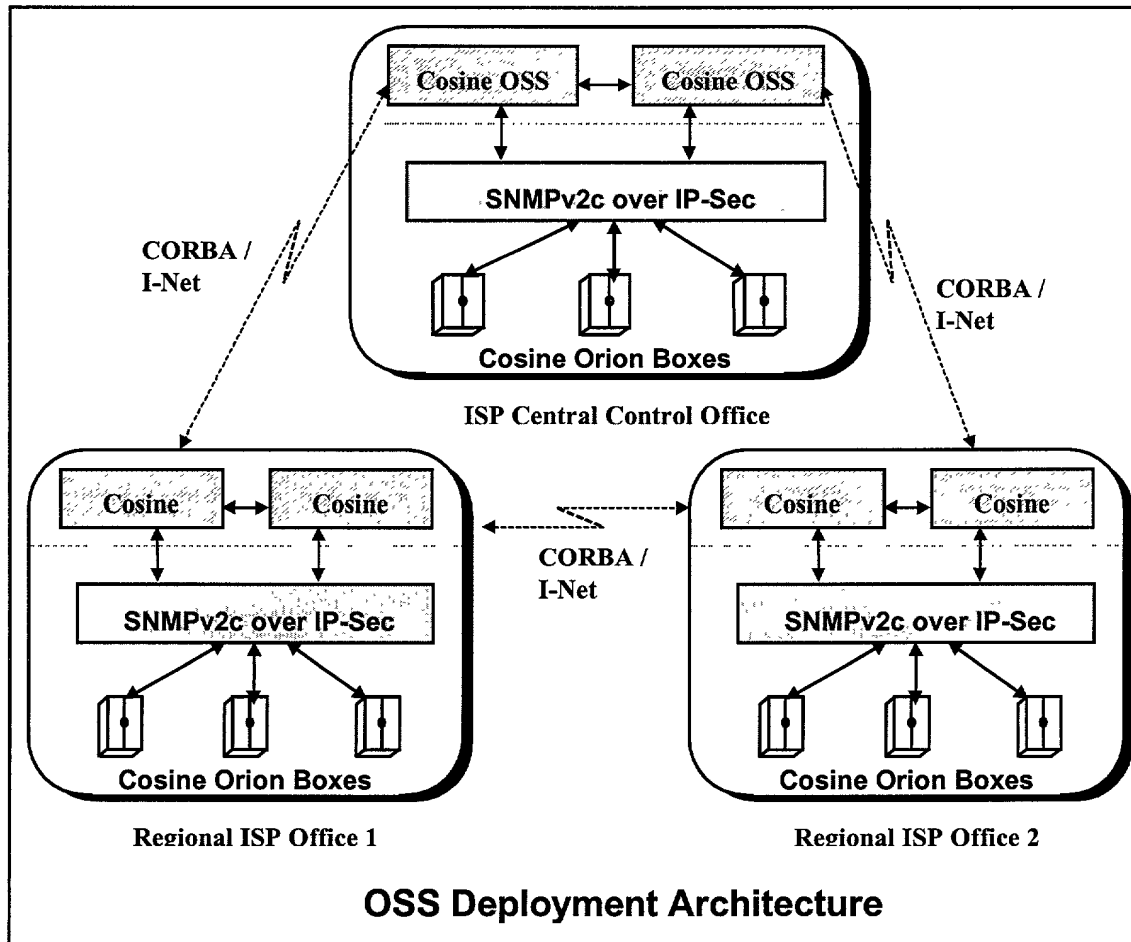
Following figure shows the logical architecture requirements of the CoSine OSS.



4.2 OSS Deployment Configuration

Following figure depicts the expected deployment scenarios for CoSine OSS. As shown in the figure below:

- A typical OSS instance will be managing multiple Orion Boxes.
- For load balancing and other purposes, a set of OSS instances could be used to manage a set of Orions.
- OSS instances can have Peer-to-Peer as well as hierarchical relationships.



4.3 External Interface Requirements

4.3.1 User Interface Requirements

OSS functionality should be accessible from a Web, JAVA, Motif and Command Line (CLI) based user interfaces. Motif interface is ~~might not be that critical but~~ support for CLI is a must. ~~At this time, it's Do not known whichat~~ protocol is / will be used by CLI programs ~~will use to access OSS~~ (a CORBA based client program could ~~have bea~~ CLI based as well).

4.3.2 Hardware Requirements

OSS server is expected to run on a Sun/Solaris box.

Appendix B

IPSX IP Flow Cache Architecture

327786

Number: CCC-NNNNN-RR

Version: 1.0

Date: August 6, 2000

Author's Name(s):

Abraham Rabindranath Mathews, Kevin Lin,
Kiho Yum, Naveed Alam, Podibanda Kurrupu

Reviewers

	Name	Check
1	Manager	✓
2	Project Lead	✓
3	Reviewer 1	✓
4	Reviewer 2	✓

Document History

[illegible]

Table of Contents

1	INTRODUCTION	1
2	BACKGROUND	1
2.1	TERMS AND ACRONYMS	1
2.2	GLOSSARY	1
3	MULTI-SERVICE CACHE DESIGN	2
4	MULTIPLE CACHE LOOKUPS FOR COMPLEX FLOWS	3
5	DISTRIBUTED FLOW CACHES	3

1 Introduction

This document is intended to describe main features of IPSX IP Flow Cache Architecture. The cache system is designed to reduce the processing overhead associated with the IP traffic passing through the system. The flow cache based forwarding helps in achieving better and more deterministic throughput and latency goals. It also helps the system resources like CPU, Memory and Interconnect Bandwidth to scale better as traffic load increases through the system.

There are three major areas identified in IFCA that have contributed the most in achieving the required performance and scalability goals.

- 1- Multi Service Forwarding Cache
- 2- Multiple Cache Lookups per Flow
- 3- Distributed Flow Cache

2 Background

There is another document by Abraham Mathews "IPSX Packet Processing Requirements" that helps in understanding the underlying problem and suggested solutions.

2.1 Terms and Acronyms

IPSec	Protocol for providing security at the IP Layer
L2TP	Layer 2 Tunneling Protocol that runs on top of UDP layer
GRE	Generic Encapsulation, a layer 4 protocol
NAT	Network Address Translation
SPF	Stateful Packet Filter
PE	Processing Elements in IPSX System
IPNOS	IP Network Operating System
VR	Virtual Router
IFCA	IP Flow Cache Architecture
IPSX	Cosine Communications IP Service Switch 9000

2.2 Glossary

L3 Flow	Only Layer 3 attributes are used to define the flow. The attributes used are <Destination IP Address, Source IP Address, Type of Service, IP Protocol>
L4 Flow	Attributes are used from layers beyond L3 to define the flow. Some protocols that contribute in L4 flows are TCP, UDP, ICMP, GRE and IPSec. In case of TCP or UDP the Layer 4 attributes used are <Destination Port, Source Port>
Packet Flow	General term used for L3 or L4 flows depending on the attribute types used.
Strict L3 Flow	flow consists of IP fragments; L4 information is not present in all fragments
IP Forwarding Services	Represent processing requirements for the ingress and egress packet flows in IPSX. The basic services consist of Routing and Forwarding services. The extended services consist

CoSine Communications Inc.	Product Requirements Document	
Company Confidential For Internal Circulation only	327786	1

of VPN and security related services like NAT, L2TP, Firewall, IPSec, SPF etc. These services are optional and can be selectively enabled by the IPSX user.

Host IPSX itself, Host addressed packets are incoming packets with destination IP address as one of the IPSX's owned IP address

Configured Topology Given a Subscriber VR and its ISP VR, consider the associated graph. This graph's nodes are the PE's executing software/ services and the edges of the graph are the communication links between the software/ services. This graph is the configured topology.

Flow Topology Based on the actual dynamic flow of L3 and L4 packets a graph of software/ services and the communication links between the software/ services can be constructed. This graph is the flow topology. The flow topology is derived from and would be a sub-graph of the configured topology if shortcuts are not employed. When shortcuts are employed, there is no direct graph theoretic correlation between the configured and flow topologies.

3 Multi-Service Cache Design

The IPSX Packet Forwarding goes through complex set of services. The basic or mandatory services consist of Routing and Forwarding services. Traditional Routers normally cache the necessary information about basic services that help in forwarding L3 flows with reduced processing overhead. If the extended services are enabled, they bring their own processing overhead in the IP forwarding path as shown in Figure 1. The extended services are normally independently designed and they bring a wide variety of processing requirements. The ability to recognize packet flows and apply flow specific processing reduces lot of processing overhead associated with these services.

The IFCA allows extended services to create their own cache contexts during the initial learning phase of the packet flows. The cache context consists of all the necessary information needed to process a service. Once the learning phase is done, the cache entry gets filled with an ordered set of operations based on the pushed cache contexts and is inserted into the hash table. The following packets belonging to the same flow find the cache entry in the hash table and go through only cache-enabled operations. The IFCA is scalable and extensible to include cacheable support for new services as they get implemented in IPSX platform.

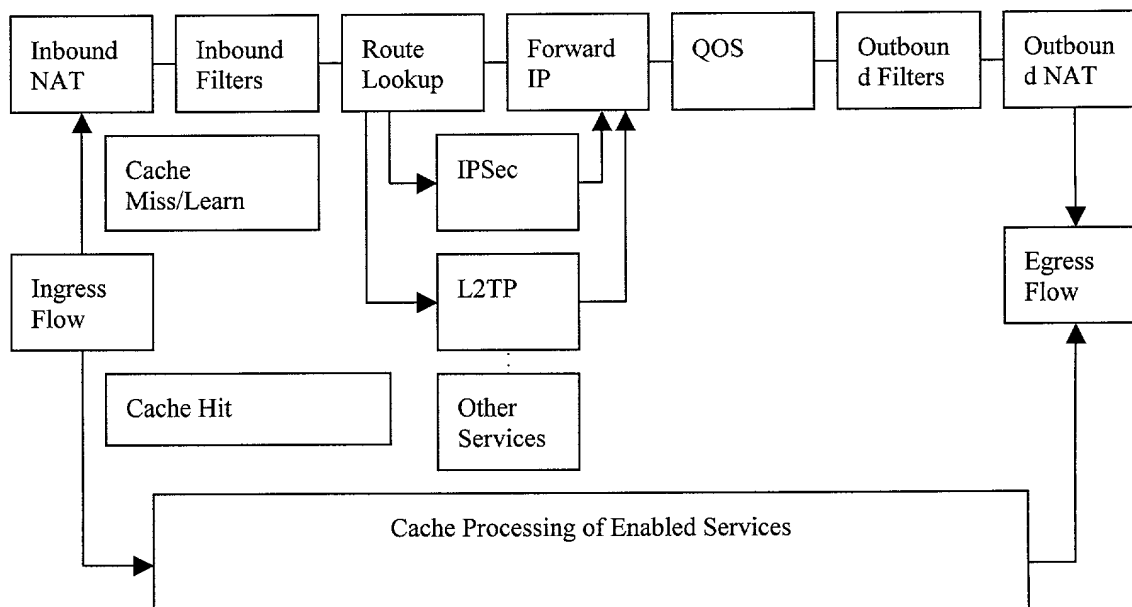


Figure 1

4 Multiple Cache Lookups for Complex Flows

The Strict L3 flows, tunneled flows and encrypted flows bring another level of complexity to the flow processing. During processing, the packet's L3 and L4 attributes transform from one type to the other. In order to keep the processing within the scope of cache operations, IFCA applies packet classifications multiple times as the L3 and L4 attributes change. This results in multiple cache lookups for the same ingress flow and enables VR to finish processing of packets through cache operations only. The traditional routers normally switch to slow path to process more complex packet flows.

As an example consider Ingress flow consisting of host addressed, encrypted, tunneled and fragmented packets. Figure 2 shows processing of Strict L3 ingress Flow of P1, P2 and P3 packets. The ingress flow goes through 3 cache lookups before turning into egress flow.

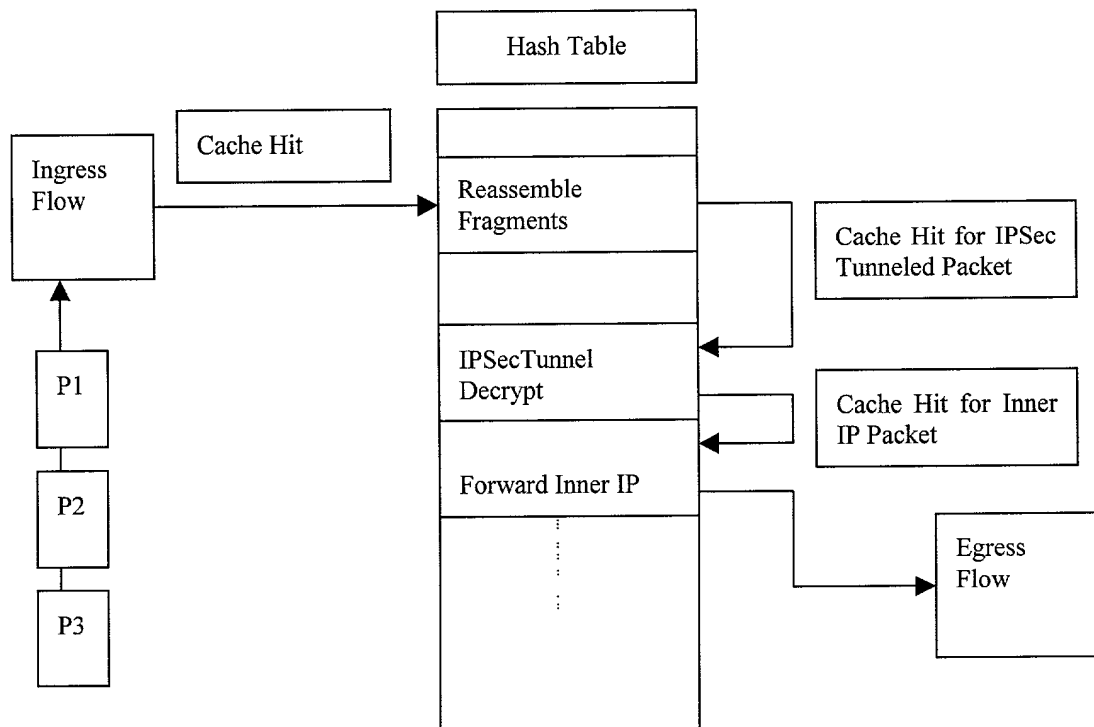


Figure 2

5 Distributed Flow Caches

IPSX provides cluster based distributed Processor Architecture. The multistage interconnect provides logical meshed topology for all PE-PE communication channels. Some PEs are equipped with special hardware like Crypto PEs to provide dedicated services. As the VR enables extended services, the IPNOS dynamically enables and links the services to the VR. Some of these services could end up on different PEs spread throughout the system. The VR cannot afford to limit its internal routing intelligence to one PE to process multiple services on different PEs. This can result into Star like flow topology putting more load on the resources of one PE doing all the internal routing for the VR.

In the process of applying services to the packet flows, the VR acquires sufficient knowledge of the internal routing of the flows. The IFCA allows the VR to take advantage of this knowledge and underlying meshed topology of PEs. The VRs can dynamically distribute the information about their flow topology to other VRs in the system. This results in optimum flow paths within the system where packet flows can directly reach the service PEs from any VR.

As an example consider two VRs, Access and ISP VR. The ISP VR is doing simple forwarding for Access VR. The Access VR has two services enabled, Intrusion Detection and IPSec. The configured topology in Figure 3 shows that Access VR has the initial knowledge of the enabled services.




Figure 3

There are two L4 Flows coming to Access VR from ISP VR. The flow F1 needs Intrusion Detection only and F2 needs Intrusion detection as well as IPSec decryption. The Access VR sends flow redirection messages to ISP VR which inserts new entries for routing of F1 and F2 flow. This redefines flow topology for both the flows as shown in figure 4

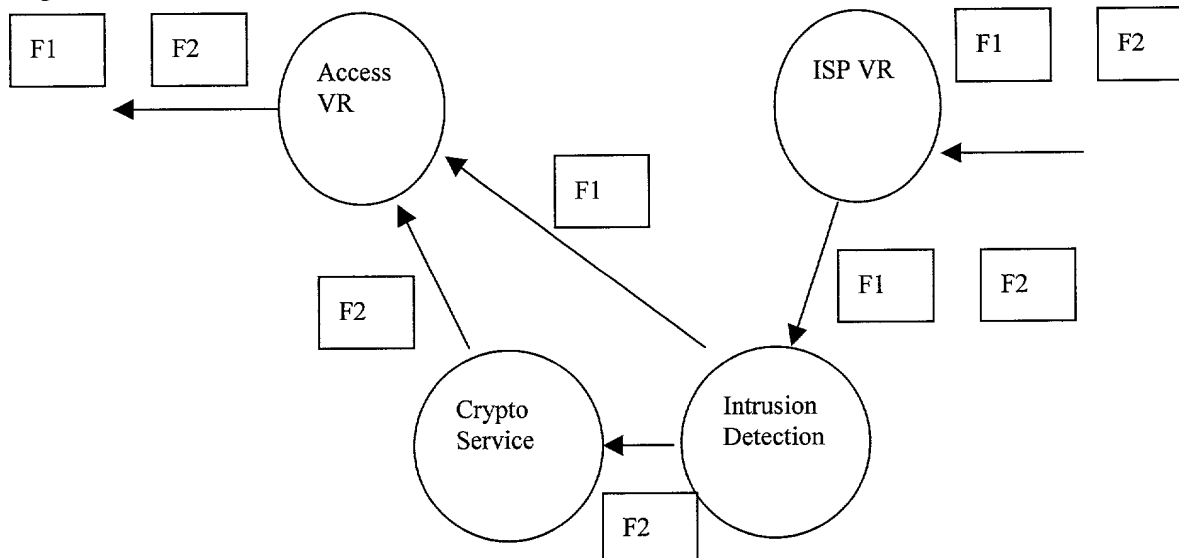



Figure 4

United States Patent & Trademark Office
Office of Initial Patent Examination -- Scanning Division



Application deficiencies were found during scanning:

☐ Page(s) _____ of _____ were not present
for scanning. (Document title)

☐ Page(s) _____ of _____ were not present
for scanning. (Document title)

☒ Scanned copy is best available. Drawings , appendix A & B